NEHRP Final Technical Report

Award No. 06HQGR0046

Model Uncertainty in Earthquake Hazard Analysis

Principal Investigator: Ralph Archuleta
Institute for Crustal Studies
University of California, Santa Barbara
Santa Barbara, CA 93106
805-893-8441
ralph@crustal.ucsb.edu

Co-PI: Jean Carlson
Department of Physics
University of California, Santa Barbara
Santa Barbara, CA 93106
805-893-8345
carlson@physics.ucsb.edu

Abstract

We use GPS data from the well-recorded 2004 M_w6.0 Parkfield Earthquake to further probe uncertainties in kinematic models. We find that the inversion for this data set is poorly resolved at depth and near the edges of the fault. In such an underdetermined inversion, it is possible to obtain spurious structure in poorly resolved areas. We demonstrate that a nonuniform grid with grid spacing matching the local resolution length on the fault outperforms small uniform grids, which generate spurious structure in poorly resolved regions, and large uniform grids, which lose recoverable information in well-resolved areas of the fault. The nonuniform grid correctly averages out large-scale structure in poorly resolved areas while recovering small-scale structure near the surface.

In addition to probing model uncertainties in earthquake source models, we also examine the effect of model uncertainty in Probabilistic Seismic Hazard Analysis (PSHA). While methods for incorporating parameter uncertainty of a particular model in PSHA are well-understood, methods for incorporating model uncertainty are more difficult to implement due to the high degree of dependence between different earthquake-recurrence models. We show that the method used by the 2002 Working Group on California Earthquake Probabilities (WGCEP-2002) to combine the probability distributions given by multiple earthquake recurrence models has several adverse effects on their result. In particular, WGCEP-2002 uses a linear combination of the models which ignores model dependence and leads to large uncertainty in the final hazard estimate. Furthermore, model weights were chosen based on data, which has the potential to systematically bias the final probability distribution. The weighting scheme used in the Working Group report also produces results that depend upon an arbitrary ordering of models. In addition to analyzing current statistical problems, we present alternative methods for rigorously incorporating model uncertainty into PSHA.

I. Model Uncertainty in Probabilistic Seismic Hazard Analysis

The goal of probabilistic seismic hazard analysis (PSHA) is to provide a quantitative estimate of the likelihood of exceeding a given threshold of earthquake-caused ground motions in a specific region during a given time period [SSHAC, 1997]. PSHA is characterized by deep uncertainty, for not only is there parameter uncertainty regarding the values of various model inputs needed to estimate hazard, there is also model uncertainty. This type of uncertainty relates to the statistical signatures for hazard, that is, how best to represent the earthquake renewal process in a recurrence model. While methods for incorporating parameter uncertainty are widely used, model uncertainty is less rigorously incorporated [Aposolakis, 1995]. Nevertheless, it is prevalent in PSHA and must be handled properly.

The 2002 Working Group on California Earthquake Probabilities (WGCEP-2002, or more concisely, WG02) differed from previous reports in that an attempt to quantify and

incorporate model uncertainty was made. Unlike previous consensus reports in 1988, 1990, and 1995 [WGCEP, 1990a, b, 1995], in which a single model was agreed upon, the WG02 report [WGCEP, 2003] used multiple models to generate the 2002 forecast. Model uncertainty was incorporated by taking a linear combination of the probability distributions given by several different models. Model uncertainty comprises a large portion of the total uncertainty in the WG02 forecast.

A. Brief Background of Previous Working Group Methodology

The 2002 Working Group on California Earthquake Probabilities (WG02) incorporated model uncertainty by taking a linear combination of the probability distributions given by several different models. While this is a laudable first step, such a methodology is strictly correct when one and only one model is correct, and the weight given that model is the probability that it is the correct model. In fact this is a very strict condition, and it is a rare case when it is satisfied. What is more likely is that each model is a different simplification of reality, so that the probability that a given model is correct is near zero. As the models are not collectively exhaustive, the probability weights will not sum to one. In particular, ignoring dependence between models leads to large uncertainties. Model uncertainty is indeed a large portion of the uncertainty in the WG02 forecast.

The Working Group on California Earthquake Probabilities (WGCEP) used five models to estimate earthquake probability in their report on earthquake hazards in the San Francisco Bay Region [*WGCEP*, 2003]. The first of these models, the Poisson Model, assumes that earthquakes randomly occur with a time-independent probability. This model has only one parameter, λ, the average rate of earthquake occurrence. The second model, the Empirical Model, is also a one-parameter Poisson-type model, but with a different value of λ. While in the Poison Model background seismicity is taken to be the long-term, historical rate of earthquake occurrence, in the Empirical Model [*Reasenberg,* 2003], post-1906 seismicity is used for λ. As recent seismicity in the San Francisco Bay Area is lower than the historical rate, the Empirical Model predicts lower probabilities than the Poisson Model.

The third model, the Brownian Passage Time (BPT) model [*Matthews et al.*, 2002; *Kagan and Knopoff*, 1987], uses two parameters to compute risk: λ and the aperiodicity of events α. The probability density for an event as a function of time t since the last event is given by

$$f_{BPT}(t) = \sqrt{\frac{1}{2\pi\lambda\alpha^2 t^3}} e^{-\frac{\lambda(t-1/\lambda)^2}{2t\alpha^2}}. \qquad (1)$$

Note that the probability density is zero at t = 0, as a new earthquake is thought unlikely until stress reaccumulates on the fault segment. When $\alpha = 0$, earthquake occurrence is periodic, and as α approaches infinity, earthquake occurrence approaches a Poisson process. The next model, the BPT-step model, is a BPT model that also incorporates the effects of stress interactions from events on nearby faults. This interaction is incorporated into the model by a "clock change" – that is, changing the value of t in Equation 1.

The final model is the Time-Predictable Model [*Shimazaki and Nakata,* 1980]. It is used only for the San Andreas Fault segment by the working group. This model uses the slip in the last earthquake, coupled with the slip rate of the fault segment, to calculate the expected time of the next rupture. The fault is expected to rupture once all of the strain released in the last earthquake has reaccumulated on the fault.

Each of the above models employs different data to estimate earthquake hazard, and each of them has different assumptions as to what controls risk. Consequently, each of these models gives a different prediction for the level of earthquake hazard in the San Francisco Bay Area. Also, as there are uncertainties in the inputs in these models, each of these models has its own associated uncertainty. The predictions and associated uncertainties for each model are shown in Figure 1a.

Given all these different models and different predictions, what is the true earthquake risk and uncertainty for the San Francisco Bay Region? This issue arises broadly in hazard analysis. Each model gives a different probability distribution. The problem is how to combing the probability distributions from different models into a single probability distribution function that best represents true earthquake risk.

Not knowing or agreeing as to which model was best or most correct, but needing to arrive at one estimate of earthquake risk, the working group performed what is essentially a weighted averaging of the models. That is, they did a Monte-Carlo sampling of all of the weighted probability distributions for each model, so that the final answer is a function of not just the mean prediction for a given model, but the entire probability distribution given by that model. This weighting is further complicated by the fact that interaction between models was allowed, that is, one model could be in place on one fault, while another model could be describing another fault for a given Monte-Carlo run. We still refer to this as an "averaging", since neglecting these interaction effects, the result of the weighted Monte Carlo sampling is a weighted average of the probability distributions produced by each model.

B. Analysis: Current Statistical Biases in PSHA

The first step in providing a rigorous methodology to incorporate model uncertainty in PSHA is a critique of previous analyses. Below we present a preliminary analysis which illustrates key issues and points towards opportunities to improve hazard estimates.

For the case of earthquake-hazard models, does averaging make sense? In fact, there is a fundamental problem assigning probabilities to different models. Averaging can be justified if one, and only one model describes the true mechanism for generating risk, and the weights used

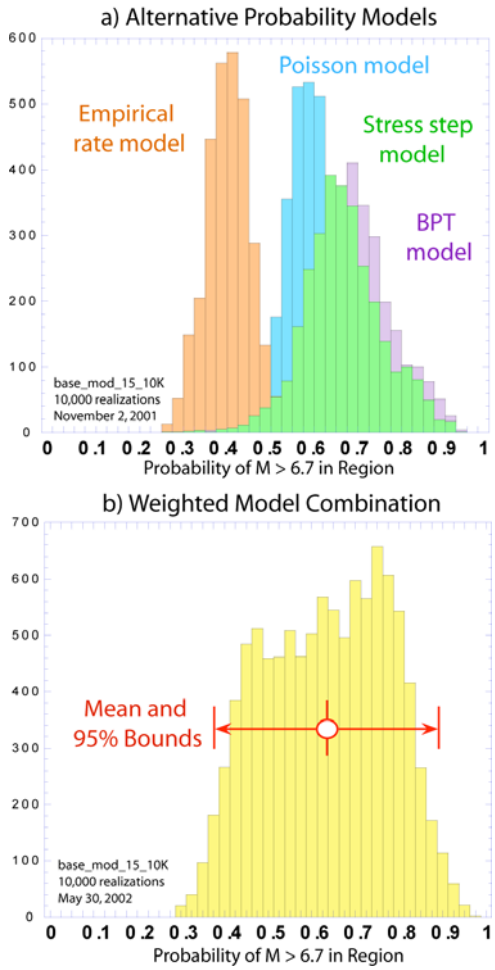

Figure 1. From the WGCEP report [2003]. Distributions of the regional probability of a M 6.7 earthquake calculated using various probability models. a) Overlapping histograms show probability calculated in 3000 iterations using each of four models separately. The shape and width of each distribution reflects epistemic uncertainty in the choice of underlying models and parameters. b) Corresponding distribution calculated in 10,000 iterations using the weighted combination of models shown in Figure 3. The broad shape of this distribution reflects the combination of distinct behaviors of the alternate models. Additional mass near P=0.8 corresponds to realizations that employ the Time-Predictable model on the San Andreas Fault, not shown in (a).

in the average reflect the probability that given model is this one "true model" [*Morgan and Henrion*, 1990]. However, this is in fact a very strict condition, and in most cases where risk is generated by the interaction of many risk factors in an unknown way, it will not be satisfied. What is more likely is that each model is a different simplification of reality, analyzing the effect of different factors in the mitigation of risk. Nonetheless averaging of probabilities produced by different models is practiced not just in the field of earthquake hazard estimation, but in climate-change studies as well [*Lempert et al.*, 2004]. Consider our case of earthquake risk. One model, for example, the Time-Predictable Model, attempts to quantify the risk associated with the slip in the last

earthquake. Another model, the Stress-Step model, makes predictions based on the stress shadows from previous earthquakes. Averaging the predictions given by these two models is not valid since it is unlikely that earthquakes are triggered only by the slip in the last earthquake or only by the stress shadows. In reality, one would expect that both of these factors are important in earthquake risk. The best prediction would use all available information to arrive at a prediction.

To further demonstrate exactly why averaging models can be incorrect, consider the following example. Suppose we are trying to assess risk for a given region. This region has "risk factor A" and "risk factor B", which could be any type of input used in a hazard model to make a prediction. For example, these "risk factors" could be time since the last earthquake, recent seismicity, strain data, etc. We also have two models for earthquake occurrence on hand: Model A, which makes predictions for earthquakes based on risk factor A, and Model B, which makes predictions based on risk factor B. Based on the historical catalog and/or physical reasoning, Model A gives a risk of α for regions having risk factor A, and thus for the region in question. Similarly, Model B assigns a risk of β to the region. Now, as neither of these models is a perfect predictor, α is not equal to β. What is our best guess of earthquake risk? Should we average α and β according to our personal opinion of the validity of the models? What can we say mathematically about the possible values for the true risk?

In fact, without knowing how risk factors A and B interact, we can place *no constraint* on what the true risk is, *even if both Model A and Model B describe how risk factors A and B affect earthquake risk perfectly*. The Venn diagrams in Figure 2 illustrate this point. It is perfectly consistent with the information given that the probability of earthquake occurrence in this region is anywhere between 0% and 100%, regardless of the values of α and β. Risk factors A and B could interact as shown in Figure 2b, so that the region has no risk, or as shown in Figure 2c, in which an earthquake is nearly guaranteed. If risk factors A and B are completely independent, risk is then given by $1-(1-\alpha)(1-\beta)$, which is greater than both α and β. Many of the plausible values for risk can not be obtained by any average!

Again, weighting models can be correct, given a strict condition: that one model, and only one, describes the true mechanism for generating risk. If this condition is in fact met (it is probably not met in earthquake hazard), we can weight each model by the probability that that model is the one correct model. However, the weights in the WG02 model are not a measure of this. In fact, one of the parameters that experts used to weight the models was "the relative amount and quality of geological data on each fault (with more and better data generally favoring more weight on recurrence models)." Similarly, Freedman and Stark [2003] remark that the weights used in an earlier report are not truly priors at all. That is, they do not reflect the experts' prior probability (initial belief) that a given model is correct. In fact, the weights the working group assigns to the various models are a function of the fault in question! Figure 3 shows the weights used for each of the seven faults in the study. Clearly, these weights do not reflect the probability that a given model is the correct simplification of reality, for we do not expect that the physics of earthquake genesis changes from fault to fault.

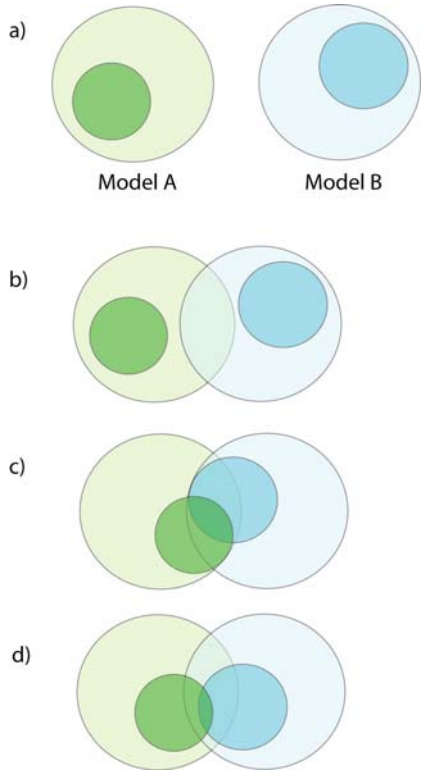

Figure 2. a) Model A calculates risk for the entire set shown having risk factor A. The smaller circle shows those members of the entire set that will in actuality have the hazardous event. Model A gives a risk of α, which is equal to the percentage of the entire set that is in the smaller circle. Similarly, Model B calculates a risk of β, which is proportional to the number of elements with Risk Factor B that are in the small blue circle. What is the risk for an element that has both Risk Factor A and Risk Factor B? It could be 0%, as shown in b), or nearly 100%, as shown in c), depending on how the risk factors interact. d) In some situations it might be wise to predict that the risk factors are independent, giving a risk of $1-(1-\alpha)\ (1-\beta)$. In each of these cases, the true risk cannot be obtained by any average of α and β, and thus, it is not correct to average the results of Model A and Model B.

Furthermore, because the weights varied from fault to fault, it was not possible to have the same model in effect on each fault for a given Monte Carlo iteration. To mitigate this problem, the models were organized in the order shown in Figure 3. Then, for a given Monte Carlo iteration, a single random number between 0 and 1 determined which model will be in effect on each fault. The method can be seen graphically from Figure 3: a given random number determines the horizontal position on the graph. A vertical line drawn at that position specifies which model is employed on each fault. For example, if the random number is 0.6, Mt. Diablo uses the Poisson model, the San Andreas uses the BPT model, and the remaining faults use the BPT-Step model. This method has several truly bizarre outcomes: certain models can interact, while others, such as the Empirical

model and the BPT model, will never both be used in the same iteration. Perhaps the most troubling result of this method, however, is that the ordering of the models changes the result. That is, the working group positioned first the Empirical Model, followed by the Poisson Model, BPT-Step Model, and BPT Model, and arranged the Time Predictable Model last, as shown in Figure 3. A different ordering of these five models, however, would lead to a different result for the combined hazard. This is troubling since the ordering is arbitrary.

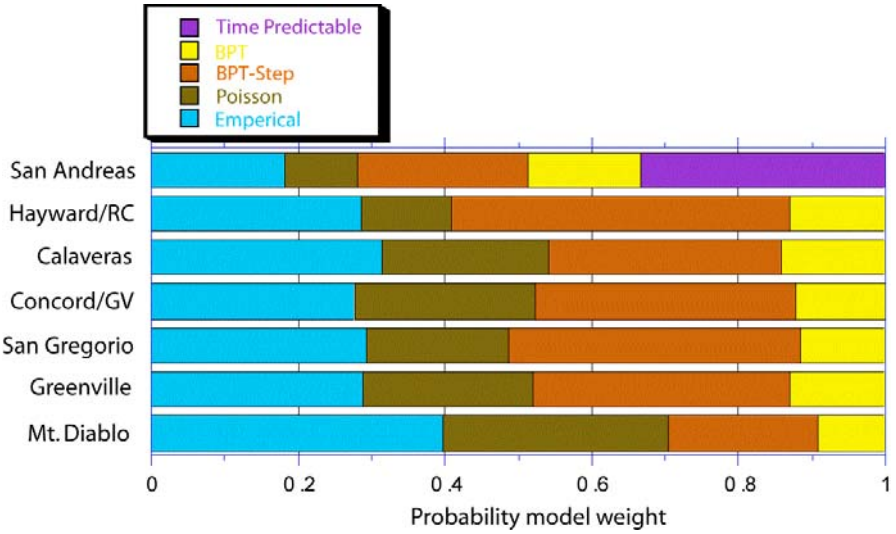

Figure 3. From the WGCEP report [2003]. Division of weight assigned to each probability model for each fault, as determined by expert opinion. The Time-Predictable model was applied only to the San Andreas fault.

At the root of this problem is the difference in model weights from fault to fault. Again, this arose because weights were based on the "relative amount and quality of geologic data" available for each fault. Weighting based on availability of data, however, can lead to systematic errors. Certainly the data available does not determine which model is the correct model, as the true physics of earthquakes is not a function of what one is able to measure. Hence it is incorrect to weight models based on this. To see how weighting based on data quality can lead to systematic errors, consider two faults, Fault A and Fault B. Suppose we also have two models of earthquake genesis, the Poisson Model and a recurrence model that gives low probabilities after earthquakes and higher probabilities after a seismic lull. Suppose Fault A has been active historically, so that data exists from past earthquakes. We might be apt to give the Poisson model a low weight for

Fault A, as the recent seismicity and abundance of data seems to make the recurrence model a better fit. After all, we should be able to make better predictions than the Poisson model for Fault A since we have such an abundance of data. Now, suppose Fault B has been seismically inactive during the historical record. We have little data, making it difficult to apply the recurrence model in this case, which requires information about past earthquakes. So we might be apt to give the Poisson Model a larger weight. But notice what we have done! The recurrence model would give a low probability for

Fault A since there were recent earthquakes and a high probability for Fault B since there were not. In each case, we weighted the model that gave the *lowest prediction* the most. In this example, giving model weights based upon data availability systematically skewed the hazard estimates to the low side.

C. Copulas

Copulas are dependence models which are ideally suited to the task of combining distributions. They are often used to combine knowledge from different experts into a single probability distribution [Clemen et al., 2000; Jouini and Clemen, 2002]. The expert-aggregation problem is similar to that of the model aggregation problem. As with model information, expert knowledge is partially dependent, since experts share knowledge. It is not a matter of which expert is right and which is wrong; rather, each expert gives additional information (presumably a function of the knowledge that differs between experts). As this is very similar to the case of model uncertainty, copulas could be used to combine multiple probability distributions from individual models into a single probability distribution.

Copulas are functions that combine univariate marginal distributions into multivariate distribution functions. A key theorem here is Sklar's Theorem [Sklar, 1959], which states that given an *n*-dimensional distribution function $H_n(x_1, ..., x_n)$ with marginal distributions $h_1(x_1), ..., h_n(x_n)$, there exists a copula C such that

$$H_n(x_1, ..., x_n) = C(h_1(x_1), ..., h_n(x_n)). \quad (2)$$

Furthermore, if $h_1(x_1), ..., h_n(x_n)$ are continuous, C is unique.

Copulas thus combine the information from the marginal distributions (in this case, the individual model probability distributions) into a single distribution H_n. Consider two models which make a prediction based on different sets of parameters. If true hazard is not a function of one set of parameters or the other, but rather an albeit complicated function of both sets of parameters, then each model is giving information that should determine the final hazard. The final probability distribution for hazard should be a function of each model output. The copula is the function that combines the two probability distributions into one distribution. From Sklar's Theorem one can see that any multivariate distribution defines a copula. The choice of copula is a function of the dependence structure of the marginals, but not a function of the marginals themselves. Choosing the copula that correctly describes the dependence between two models is the most important part of this formulation. Clemen and Reilly [1999] discuss how to use expert opinion to this end. If the models are exchangeable in terms of their dependence, an Archimedian copula is ideal, as it treats the marginal distributions symmetrically [Jouini and Clemen, 2002]. For more flexibility, the multivariate normal copula can be used [Clemen and Reilly, 1999]. It encodes dependence using pairwise correlation coefficients, for example, Spearman's ρ or Kendall's τ [Lehmann, 1966]. This can be ideal in situations where there is little data, for the statistical measures of dependence ρ or τ can be assessed with expert opinion.

D. An Example Applying Copulas to Earthquake Recurrence Models

Jouini and Clemen [2002] present a simple method using copulas in a Bayesian framework to combine multiple probability distributions from experts. We follow their method here to combine two distributions from the WG02 report: the probability distributions given by the Poisson and Empirical models. Jouini and Clemen [2002] use a uniform prior distribution for the application of Bayes' Rule. To calculate the likelihood function, they combine the individual expert probability distributions $f_i(\theta)$ using a family of copulas described by Frank [1979]. This family of copulas is indexed by a single parameter, Kendall's τ [Kendall , 1938]. For two independent and identically distributed pairs of random variables, Kendall's τ is defined as the probability of concordance minus the probability of discordance. The copula family of Frank [1979] can capture the full range of positive dependence from $\tau = 0$ (independence) to $\tau = 1$ (perfect positive dependence). Furthermore, Jouini and Clemen assume for simplicity that given a random median M_i given by the ith marginal, the marginal distributions are only dependent through their estimation errors $\theta - M_i$. Under these assumptions, the posterior probability distribution is proportional to

$$c(1 - F_1(\theta), ..., 1 - F_n(\theta))\, f_1(\theta) \ldots f_n(\theta), \qquad (3)$$

where c is the copula density function, and F_i is the cumulative probability distribution generated from the ith marginal f_i.

An example of this copula-based method is applied to the Poisson and Empirical probability distributions from the WG02 report and shown in Figure 4. We fit the individual probability distributions from Figure 1a with Gaussian distributions. Recall that each statistical model produces a distribution of values for the probability of an earthquake of magnitude 6.7 of greater. The width of the Gaussian that we fit is determined by the parameter uncertainty of the individual models. In Figure 4a, we show the regional probability distributions given by the Poisson and Empirical Models, along with an equally weighted average of the two in black. This is analogous to the WG02 methodology. Figures 4b, c, and d show examples of the copula-based aggregation for three different levels of dependence. Figure 4b shows an aggregation where the two models are assumed to be independent ($\tau = 0$). In this case, Frank's copula reduces to the independence copula $C = f_1 f_2$. If the two models are independent, then together they provide the most information, and the variance in the combined distribution is as small as possible. Higher levels of dependence result in more variance, as more weight is added to regions where the two distributions differ. Large values of Kendall's τ result in a bimodal combined distribution (Figure 4d). In the case of two Gaussian marginals with equal variance, the combined distribution would be symmetric. In general, Frank's copulas give more weight to the marginal with less variance, and thus the Empirical model is given more weight in the aggregated distributions in this example.

The copula-based aggregations shown in Figure 4b, c, and d do not assume that the Poisson and Empirical models are mutually exclusive. Rather, they treat both models as marginal distributions. The Poisson Model is a marginal distribution for historical seismicity and the Empirical Model is a marginal distribution for post-1906 seismicity.

The copula is the function that combines the two marginals into the bivariate probability density.

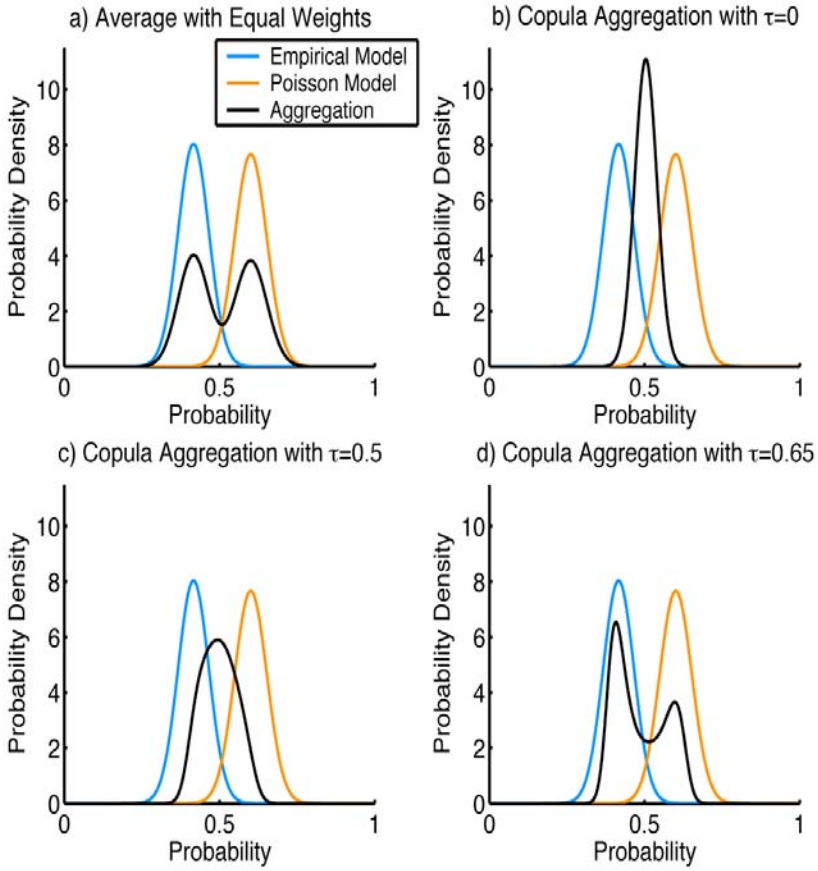

Figure 4. Combining the Poisson and Empirical Model with an equally weighted linear combination (a) leads to quite different results than a using a copula-based aggregation method (b-d). If the two models are independent, Kendall's _ equals zero and our copula reduces to the independence copula. This aggregation has the least variance (b). If we assume more dependence between the input models, then there is less total information. As we expect, in that case the combined model has higher variance, as shown (c and d).

The above approach can easily be generalized to incorporate more than two models. For each pair of models, one pairwise correlation coefficient is needed to completely define the copula from this particular family. Expert opinion could easily be used to this end. The main drawback of this formulation is that the answer is highly dependent on the type of dependence structure between the models, and thus on the copula chosen. Archimedian copulas, which treat the dependence between the marginal distributions symmetrically, are certainly the easiest to implement. The copulas we use here from Frank [1979] are members of this class. Modeling complex dependence structures,

however, requires a more sophisticated analysis. To this end, Clemen and Reilly [1999] discuss methods using the copula underlying the multivariate normal distribution. In addition, MacKenzie [1994] develops a class of copulas with even more flexibility.

Treating the model weights as probabilities (as a linear combination does) is problematic because the weights are not mutually exclusive or collectively exhaustive. Copulas provide a way to combine multiple models without abandoning probabilism.

II. Uncertainty in Earthquake Source Models

A. Resolution of GPS Data from the 2004 M6.0 Parkfield Earthquake

Given a linear inverse problem, the resolution matrix $R=G^{-g}G$ is a function of the Green's function G and the generalized inverse G^{-g}. If the inversion has perfect resolution, the resolution matrix will equal the identity matrix. In practice, the rows of R give weighted averages of the model parameters [*Menke*, 1989]. The extent to which the weights along the diagonal elements of R "leak" into neighboring elements in each row gives a measure of the resolution of each model parameter.

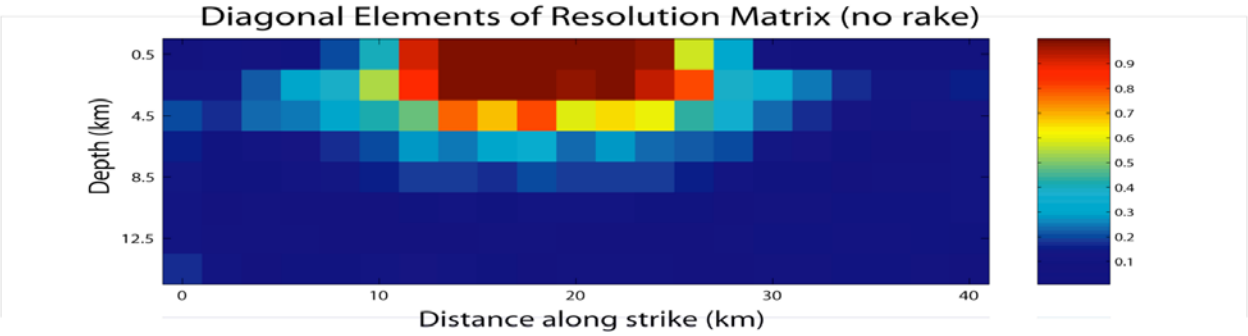

Figure 5. Diagonal elements of the resolution matrix for the GPS data from the 2004 Mw6.0 Parkfield Earthquake plotted on their corresponding subfaults. Resolution is poor at depth and near the edges of the fault.

We found that even for a well-recorded earthquake such as the 2004 M6.0 Parkfield earthquake, which was recorded by 13 near-field GPS stations, static GPS inversions are poorly resolved at depth and near the edges of the fault. Figure 5 shows the on-diagonal elements of the resolution matrix mapped onto the fault plane. In poorly resolved areas of the fault, slip is poorly constrained spatially, which can lead to artifacts at depth. The location of these artifacts, which often look very similar to asperities, is a function of the station locations and velocity structure.

B. Inversion of Parkfield GPS Data

Formulating the inverse problem in a way that is severely underdetermined can lead to spurious structure in the final model. In the inversion of Parkfield GPS data, the resolution is highly spatially variable, with a much smaller resolution length near the top and center of the fault plane. We can improve the model resolution by making the subfaults larger in poorly resolved areas. A nonuniform grid with subfaults (grid cells) that match the local resolution length on the fault plane simultaneously maximizes the recoverable information in well-resolved areas of the fault while avoiding spurious structure in poorly resolved areas.

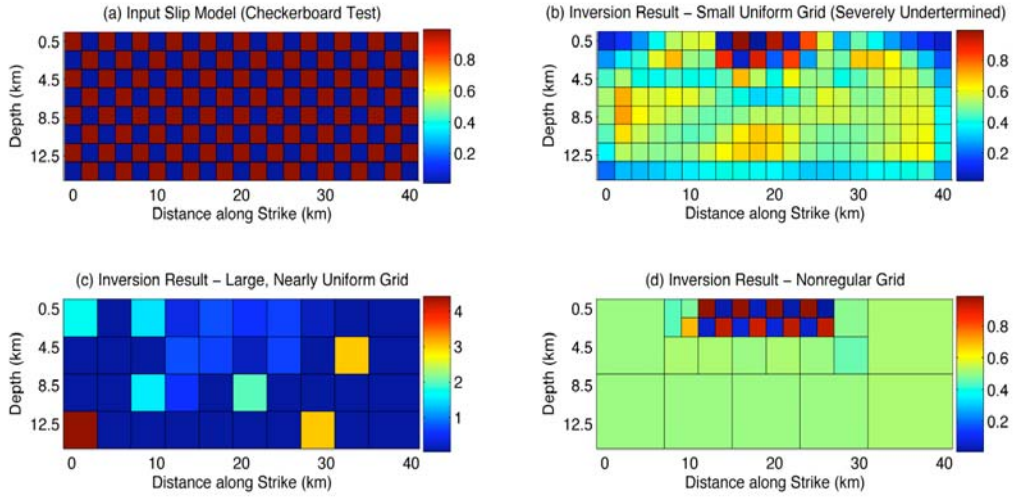

Figure 6. We generated data from a synthetic slip model (a checkerboard test) shown in (a) and inverted the data onto three different grids. b) With a small, uniform grid, spurious structure is generated at depth. c) With a larger uniform gridding, structure near the surface is lost and spurious structure is again generated at depth in part because the large subfaults near the surface are removing structure that is within the resolution length of the problem. d) With a nonuniform grid with spacing that approximates the local resolution length on the fault plane, structure is adequately recovered in well-resolved portions of the fault and spurious slip is avoided in poorly resolved areas.

We performed a traditional inversion of the Parkfield GPS data on a uniform grid and compared this result to an inversion performed on the nonuniform grid shown in Figure 6d. Figures 7a and 7b show the slip model from the uniform-grid inversion and the associated perturbation error from a Monte Carlo sampling of the errors in the GPS data. The inversion on the nonuniform grid, as shown in Figure 7c, captures the resolution error in the gridding of the fault plane. In our view, the nonuniform-grid inversion of the Parkfield GPS data is superior because it assesses both resolution and perturbation errors. In addition, it is less likely to contain artifacts because the larger subfaults at depth limit the number of free parameters.

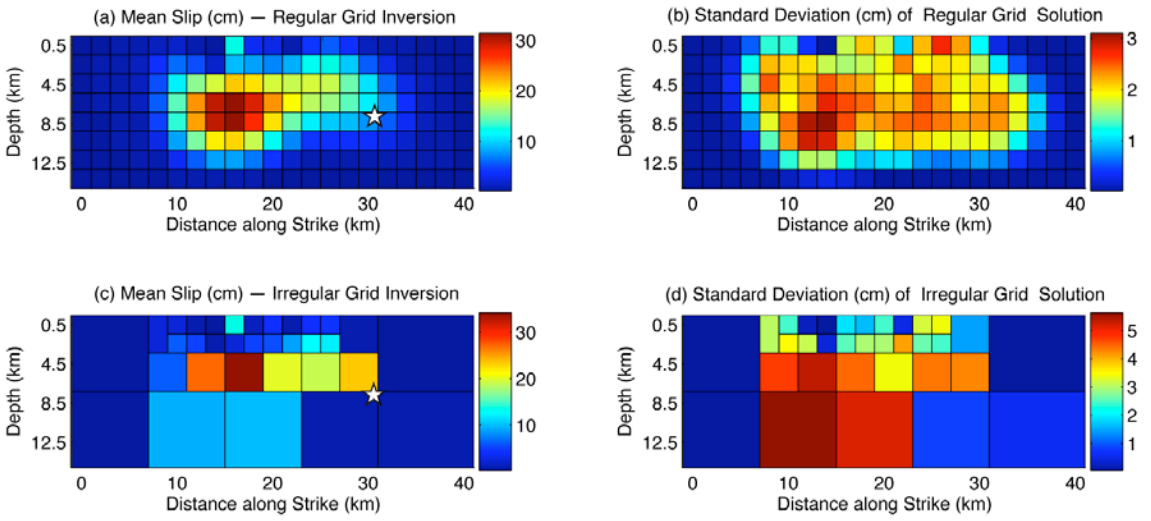

Figure 7. Inversion of Parkfield GPS data on a regular grid (a), and on an irregular grid (c), with associated perturbation errors found via Monte Carlo sampling of GPS errors (b and d). Both inversions give similar fits to the data with a variance reduction of 89-90%.

C. Two-Step Inversion with Strong-Motion Data

We constrained the final slip in an inversion of strong-motion data to match the final slip given by the GPS data, within the error bounds. The GPS inversion of the regular grid does not place enough slip near the hypocenter to satisfy strong-motion stations to the southeast. One might think that this suggests an inconsistency between the GPS and strong-motion data, but in fact the GPS inversion on the irregular grid is consistent with the strong motion data. The non-uniform grid allows us to capture resolution error, and when this is taken into account, the two datasets agree. This confirms our view that the nonuniform grid produces more reliable results with fewer artifacts. The results of the two-step inversion are shown in Figure 8.

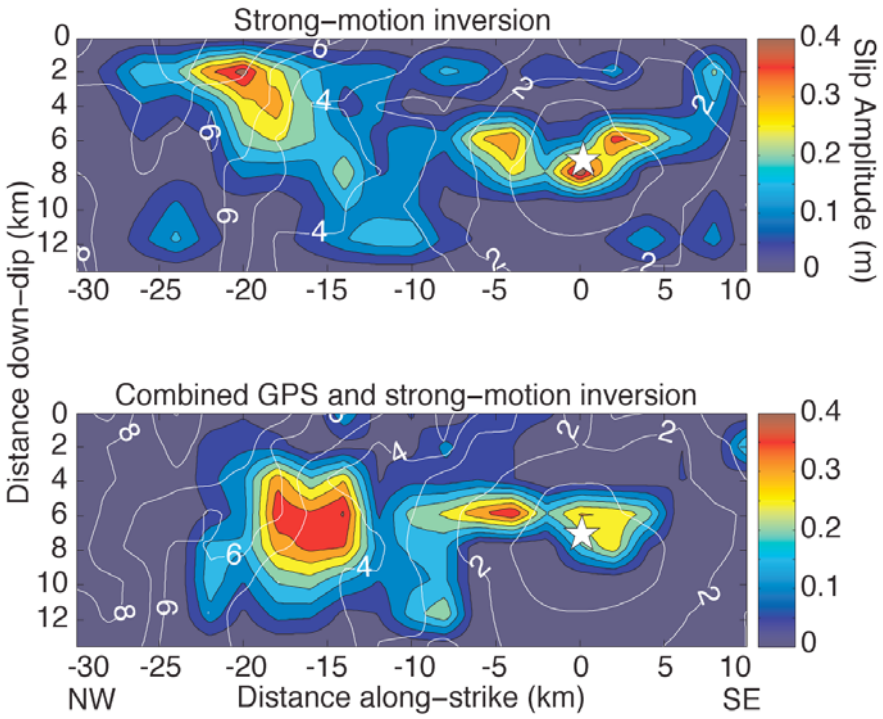

Figure 8. Inversion of strong-motion GPS data without any static-field constraint [*Liu et al.,* 2006] (a), and with the addition of a final-slip constraint derived from the GPS data inverted on a nonuniform grid (b). The addition of GPS data results in a more compact final slip distribution.

D. Resolution of Strong-Motion Data

Strong-motion data has the ability to see slip deeper than static data because the static field decays faster with source-to-receiver distance [*Aki and Richards*, 2002]. Synthetic tests of different data types confirm this (e.g., [*Delouis et al.*, 2002]). However, strong-motion data is further complicated by rupture time, which adds a strong nonlinearity to the inverse problem. We linearized the strong-motion inversion about the fiinal rupture times determined by a nonlinear simulated annealing algorithm in order to investigate the resolution and covariance of the model.

Model vectors that correspond to small singular values are unstable to data perturbations. In Figure 9 we plot the distribution of singular values for the GPS inversion and for the linearized strong-motion inversion. Even though this linearized version of the strong-motion inversion is overdetermined, many of the model parameters (particularly at depth) are unstable, and thus very sensitive to data perturbations.

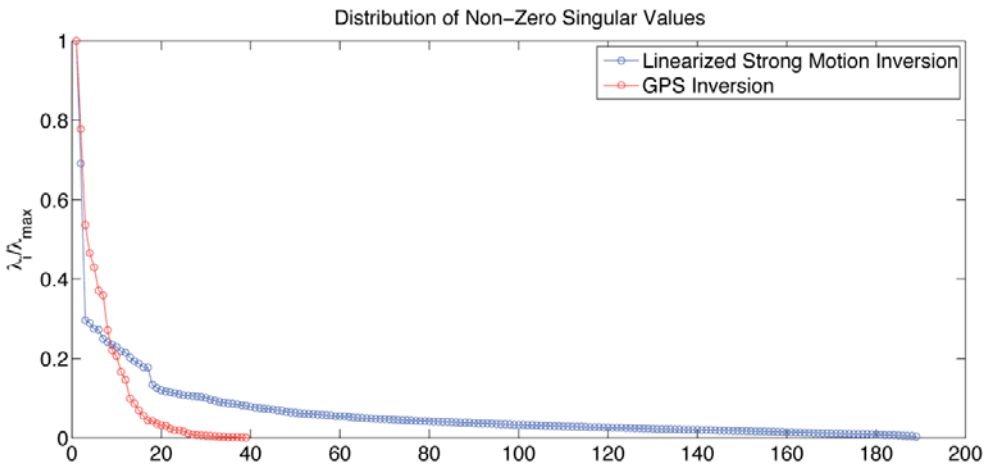

Figure 9. Distribution of singular values for the GPS inversion (red) and for a linearized strong-motion inversion (blue). The strong-motion inversion is not technically underdetermined, although many of the singular values are near zero, which means some slip vectors are unstable.

The log of the model covariance matrix, shown in Figure 10, shows that parameters at depth are far more sensitive to data perturbations than shallower parameters. This is important because a major source of perturbation error in strong-motion inversions is due to Green's function errors. Strong-motion inversions of the dynamic wavefield are more sensitive to the velocity structure than are static inversions [*Wald and Graves*, 2001]. Errors in the Green's function due to incorrect velocity structure or fault location are highly nonlinear, and can change the final slip model significantly [*Das and Suhadolc*, 1996; *Sekiguchi et al.*, 2000]. A thorough quantification of errors in kinematic inversions will allow for the determination of robust features in the models, which will allow researchers to draw firmer conclusions from this information.

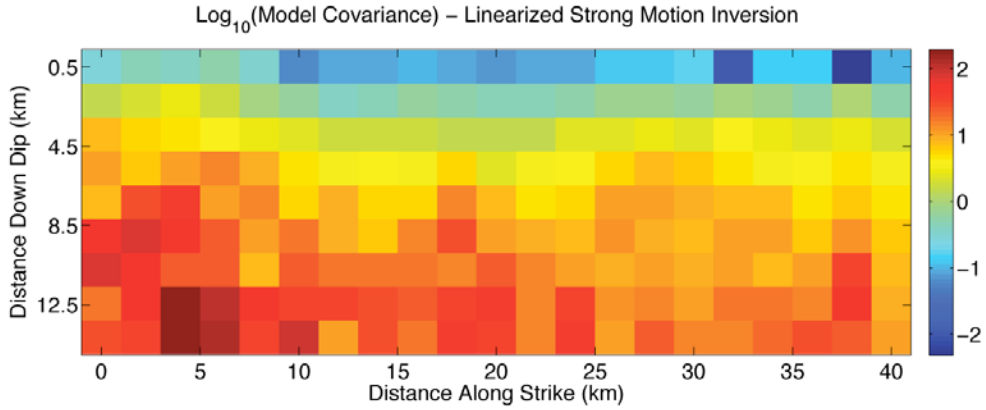

Figure 10. The model covariance, which shows how sensitive model parameters are to perturbations in the data, is shown plotted onto the fault plane for a linearized strong-motion inversion.

Publications:

Page, M. T. and J. M. Carlson (2006), Methodologies for Earthquake Hazard Assessment: Model Uncertainty and the WGCEP-2002 Forecast, *Bull. Seism. Soc. Am., 96, 5, doi: 10.1785/0120050195.*

Page, Morgan, Susana Custódio, Ralph J. Archuleta, and J. M. Carlson. Constraining Earthquake Source Inversions with GPS Data 1: Resolution Based Removal of Artifacts, submitted to *J. Geophys. Res. - Solid Earth.*

Custódio, Susana, Morgan T. Page, Ralph J. Archuleta, and J.M. Carlson. Constraining Earthquake Source Inversions with GPS Data 2: A Two-Step Approach to Combine Seismic and Geodetic Datasets, submitted to *J. Geophys. Res. - Solid Earth.*

References:

Aki, K., and P. G. Richards (2002), Quantitative Seismology, second ed., *University Science,* Sausalito, California.

Aposolakis, G. (1995), Model Uncertainty: Its Characterization and Quantification, chap. A Commentary On Model Uncertainty, pp. 13–22, *Center for Reliability Engineering*, University of Maryland.

Clemen, R. T., and T. Reilly (1999), Correlations and copulas for decision and risk analysis, *Management Science*, 45, 208–224.

Clemen, R. T., G. W. Fischer, and R. L. Winkler (2000), Assessing dependence: Some experimental results, *Management Science*, 46, 1100–1115.

Das, S., and P. Suhadolc (1996), On the inverse problem for earthquake rupture: The Haskell-type source mode, *J. Geophys. Res.,* 101(B3), 5725–5738.

Delouis, B., D. Giardini, P. Lundgren, and J. Salichon (2002), Joint inversion of InSAR, GPS, teleseismic, and strong-motion data for the spatial and temporal distribution of earthquake slip: Application to the 1999 Izmit mainshock, *Bull. Seismol. Soc. Am.*, 92(1), 278–299.

Frank, M. J. (1979), On the simultaneous associativity of $F(x, y)$ and $x + y - F(x, y)$, *Aeq. Math.,* 19, 194–226.

Freedman, D. A., and P. B. Stark (2003), What is the chance of an earthquake?, *NATO Science Series IV: Earth and Environmental Sciences*, 32, 201–213.

Jouini, M. N., and R. T. Clemen (2002), Copula models for aggregating expert opinions, *Operations Research*, 44, 444–457.

Kagan, Y. Y., and L. Knopoff (1987), Random stress and earthquake statistics: Time dependence, *Geophys. J. R. Astr. Soc.*, 88, 723–731.

Kendall, M. G. (1938), A new measure of rank correlation, *Biometrika*, 30, 81–93.

Lehmann, E. L. (1966), Some concepts of dependence, *Ann. Math. Statist.*, 37, 1137–1153.

Lempert, R., N. Nakicenovic, D. Sarewitz, and M. Schlesinger (2004), Characterizing climate-change uncertainties for decision-makers, *Climatic Change*, 65, 1–9.

Liu, P., S. Custódio, and R. J. Archuleta (2006), Kinematic inversion of the 2004 M_w6.0 Parkfield earthquake including an approximation to site effects, *Bull. Seismol. Soc. Am.*, **96**, 4B, doi: 10.1785/0120050826

MacKenzie, G. (1994), Approximately maximum-entropy multivariate distributions with specified marginals and pairwise correlations, *Ph.D. thesis*, University of Oregon.

Matthews, M. V., W. L. Ellsworth, and P. A. Reasenberg (2002), A Brownian model for recurrent earthquakes, *Bull. Seismol. Soc. Am.*, 92, 2233–2250.

Menke, W. (1989), Geophysical Data Analysis: Discrete Inverse Theory, *Academic Press,* San Diego.

Morgan, M. G., and M. Henrion (1990), Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis, 67-68 pp., *Cambridge University Press*, Cambridge.

Reasenberg, P. A., T. C. Hanks, and W. H. Bakun (2003), An empirical model for earthquake probabilities in the San Francisco Bay Region, California: 2002-2031, *Bull. Seismol. Soc. Am.*, 93, 1-13.

Sekiguchi, H., K. Irikura, and T. Iwata (2000), Fault geometry at the rupture termination of the 1995 Hyogo-ken Nanbu earthquake, *Bull. Seismol. Soc. Am.*, 90, 117–133.

Shimazaki, K., and T. Nakata (1980), Time-predictable recurrence model for large earthquakes, *Geophysical Research Letters*, 7, 279–282.

Sklar, A. (1959), Fonctions de répartition à n dimentions et leurs marges, *Pub. Inst. Stat. Univ. Paris*, 8, 229–231.

SSHAC (1997), Recommendations for probabilistic seismic hazard analysis: Guidance on uncertainty and use of experts, Tech. Rep. NUREG/CR-6372, *Lawrence Livermore National Laboratory*.

Wald, D. J., and R. W. Graves (2001), Resolution analysis of finite fault source inversion using one- and three-dimensional Green's functions: 2. Combining seismic and geodetic data, *J. Geophys. Res.*, 106, B5, 8767–8788.

WGCEP (1990a), Probabilities of large earthquakes occurring in California on the San Andreas fault, *Tech. Rep. OFR 88-398*, USGS.

WGCEP (1990b), Probabilities of large earthquakes in the San Francisco Bay Region, California, *U.S. Geological Survey Circular*, p. 51.

WGCEP (1995), Seismic hazards in southern California: probable earthquakes, 1994-2024, *Bull. Seismol. Soc. Am.*, 85, 379–439.

WGCEP (2003), Earthquake probabilities in the San Francisco Bay Region: 2002-2031, *Tech. Rep. OFR 03-214*, USGS.